arXiv:1605.06439v1 [cs.LG] 20 May 2016

Combining Adversarial Guarantees and Stochastic Fast Rates in Online Learning

Wouter M. Koolen
Centrum Wiskunde & Informatica
wmkoolen@cwi.nl

Peter Grünwald
Centrum Wiskunde & Informatica
and Leiden University
pdg@cwi.nl

Tim van Erven
Leiden University
tim@timvanerven.nl

May 23, 2016

Abstract

We consider online learning algorithms that guarantee worst-case regret rates in adversarial environments (so they can be deployed safely and will perform robustly), yet adapt optimally to favorable stochastic environments (so they will perform well in a variety of settings of practical importance). We quantify the friendliness of stochastic environments by means of the well-known Bernstein (a.k.a. generalized Tsybakov margin) condition. For two recent algorithms (Squint for the Hedge setting and MetaGrad for online convex optimization) we show that the particular form of their data-dependent individual-sequence regret guarantees implies that they adapt automatically to the Bernstein parameters of the stochastic environment. We prove that these algorithms attain fast rates in their respective settings both in expectation and with high probability.

1 Introduction

We consider online sequential decision problems. We focus on full information settings, encompassing such interaction protocols as online prediction, classification and regression, prediction with expert advice or the Hedge setting, and online convex optimization (see Cesa-Bianchi and Lugosi 2006). The goal of the learner is to choose a sequence of actions with small regret, i.e. such that his cumulative loss is not much larger than the loss of the best fixed action in hindsight. This has to hold even in the worst case, where the environment is controlled by an adversary. After three decades of research there exist many algorithms and analysis techniques for a variety of such settings. For many

settings, adversarial regret lower bounds of order $\sqrt{T}$ are known, along with matching individual sequence algorithms [Shalev-Shwartz, 2011].

A more recent line of development is to design adaptive algorithms with regret guarantees that scale with some more refined measure of the complexity of the problem. For the Hedge setting, results of this type have been obtained, amongst others, by Cesa-Bianchi et al. [2007], De Rooij et al. [2014], Gaillard et al. [2014], Sani et al. [2014], Koolen et al. [2014], Koolen and Van Erven [2015], Luo and Schapire [2015], Wintenberger [2015]. Interestingly, the price for such adaptivity (i.e. the worsening of the worst-case regret bound) is typically extremely small (i.e. a constant factor in the regret bound). For *online convex optimization* (OCO), many different types of adaptivity have been explored, including by [Crammer et al., 2009, Duchi et al., 2011, McMahan and Streeter, 2010, Hazan and Kale, 2010, Chiang et al., 2012, Steinhardt and Liang, 2014, Koolen and Van Erven, 2016].

Here we are interested in the question of whether such adaptive results are strong enough to lead to improved rates in the stochastic case when the data follow a "friendly" distribution. In specific cases it has been shown that fancy guarantees do imply significantly reduced regret. For example Gaillard et al. [2014] present a generic argument showing that a certain kind of second-order regret guarantees implies constant expected regret (the fastest possible rate) for i.i.d. losses drawn from a distribution with a gap (between expected loss of the best and all other actions). In this paper we significantly extend this result. We show that a variety of individual-sequence second-order regret guarantees imply fast regret rates for distributions under much milder stochastic assumptions. In particular, we will look at the Bernstein condition (see Bartlett and Mendelson 2006), which is the key to fast rates in the batch setting. This condition provides a parametrised interpolation (expressed in terms of the Bernstein exponent $\kappa \in [0,1]$) between the friendly gap case ($\kappa = 1$) and the stochastic worst case ($\kappa = 0$). We show that appropriate second-order guarantees automatically lead to adaptation to these parameters, for both the Hedge setting and for OCO. In the Hedge setting, we build on the guarantees available for the Squint algorithm [Koolen and Van Erven, 2015] and for OCO we rely on guarantees achieved by MetaGrad [Koolen and Van Erven, 2016] to obtain regret rates of order $T^{\frac{1-\kappa}{2-\kappa}}$ (Theorem 3). We show this, not just in expectation, but also with high probability. Our proofs use that, for bounded losses, the Bernstein condition is equivalent to the so-called Central condition [Van Erven et al., 2015], which provides control over a martingale-type quantity that captures the second-order part of the bounds (Lemma 8). The rates we obtain include the slow worst-case $\sqrt{T}$ regime for $\kappa = 0$ and the fastest (doubly) logarithmic regime for $\kappa = 1$.

The next section introduces the two settings we consider and the individual sequence guarantees we will use in each. It also reviews the stochastic criteria for fast rates and presents our main result. In Section 3 we consider a variety of examples illustrating the breadth of cases that we cover. In Section 4 we prove that second-order guarantees imply adaptation to Bernstein conditions.

2 Setup

2.1 Hedge Setting

We start with arguably the simplest setting of online prediction, the Hedge setting popularized by Freund and Schapire [1997]. To be able to illustrate the full reach of our stochastic assumption we will use a minor extension to countably infinitely many actions $k \in \mathbb{N} = \{1, 2, \ldots\}$, customarily called experts. The protocol is as follows. Each round t the learner plays a probability mass function $w_t = (w_t^1, w_t^2, \ldots)$ on experts. Then the environment reveals the losses $\ell_t = (\ell_t^1, \ell_t^2, \ldots)$ of the experts, where each $\ell_t^k \in [0, 1]$. The learner incurs loss $\langle w_t, \ell_t \rangle = \sum_k w_t^k \ell_t^k$. The regret after T rounds compared to expert k is given by

$$R_T^k \coloneqq \sum_{t=1}^{T} \left(\langle w_t, \ell_t \rangle - \ell_t^k \right).$$

The goal of the learner is to keep the regret small compared to any expert k. We will make use of *Squint* by Koolen and Van Erven [2015], a self-tuning algorithm for playing w_t. Koolen and Van Erven [2015, Theorem 4] show that Squint with prior probability mass function $\pi = (\pi^1, \pi^2, \ldots)$ guarantees

$$R_T^k \leq \sqrt{V_T^k K_T^k} + K_T^k \quad \text{where} \quad K_T^k = O(-\ln \pi^k + \ln \ln T) \qquad \text{for any expert } k. \tag{1}$$

Here $V_T^k \coloneqq \sum_{t=1}^{T} \left(\langle w_t, \ell_t \rangle - \ell_t^k \right)^2$ is a second-order term that depends on the algorithm's own predictions w_t. It is well-known that with K experts the worst-case lower bound is $\Theta(\sqrt{T \ln K})$ [Cesa-Bianchi and Lugosi, 2006, Theorem 3.7]. Taking a fat-tailed prior, for example $\pi^k = \frac{1}{k(k+1)}$, and using $V_T^k \leq T$, the above bound implies $R_T^k \leq O\left(\sqrt{T(\ln k + \ln \ln T)}\right)$, matching the lower bound in some sense for all k simultaneously.

The question we study in this paper is what becomes of the regret when the sequence of losses $\ell_1, \ell_2, \ldots$ is drawn from some distribution $\mathbb{P}$, not necessarily i.i.d. But before we expand on such stochastic cases, let us first introduce another setting.

2.2 Online Convex Optimization (OCO)

We now turn to our second setting called *online convex optimization* [Shalev-Shwartz, 2011]. Here the set of actions is a compact convex set $\mathcal{U} \subseteq \mathbb{R}^d$. Each round t the learner plays a point $w_t \in \mathcal{U}$. Then the environment reveals a convex loss function $\ell_t : \mathcal{U} \to \mathbb{R}$. The loss of the learner is $\ell_t(w_t)$. The regret after T rounds compared to $u \in \mathcal{U}$ is given by

$$R_T^u \coloneqq \sum_{t=1}^{T} \left(\ell_t(w_t) - \ell_t(u) \right).$$

The goal is small regret compared to any point $u \in \mathcal{U}$. A common tool in the analysis of algorithms is the linear upper bound on the regret obtained from

convexity of ℓ_t (at non-differentiable points we may take any sub-gradient)

$$R_T^u \;\le\; \tilde{R}_T^u \;\coloneqq\; \sum_{t=1}^{T} \langle w_t - u, \nabla \ell_t(w_t) \rangle.$$

We will make use of (the full matrix version of) *MetaGrad* by Koolen and Van Erven [2016]. In their Theorem 8, they show that, simultaneously, $\tilde{R}_T^u \le O\left(DG\sqrt{T}\right)$ and

$$\tilde{R}_T^u \;\le\; \sqrt{V_T^u K_T} + DGK_T \quad \text{where} \quad K_T \;=\; O(d \ln T) \qquad \text{for any } u \in \mathcal{U}, \tag{2}$$

where D bounds the two-norm diameter of $\mathcal{U}$, G bounds $\|\nabla \ell_t(w_t)\|_2$ the two-norm of the gradients and $V_T^u \coloneqq \sum_{t=1}^{T} \langle w_t - u, \nabla \ell_t(w_t) \rangle^2$. The first bound matches the worst-case lower bound. The second bound (2) may be a factor $\sqrt{K_T}$ worse, as $V_T^u \le G^2 D^2 T$ by Cauchy-Schwarz. Yet in this paper we will show fast rates in certain stochastic settings arising from (2). To simplify notation we will assume from now on that $DG = 1$ (this can always be achieved by scaling the loss).

To talk about stochastic settings we will assume that the sequence ℓ_t of loss functions (and hence the gradients $\nabla \ell_t(w_t)$) are drawn from a distribution $\mathbb{P}$, not necessarily i.i.d. This includes the common case of linear regression and classification where $\ell_t(u) = \text{loss}(\langle u, x_t \rangle, y_t)$ with (x_t, y_t) sampled i.i.d. from some distribution and loss a fixed one-dimensional convex loss function (e.g. square loss, absolute loss, log loss, hinge loss, ...).

2.3 Parameterized Family of Stochastic Assumptions

We now recall the Bernstein [Bartlett and Mendelson, 2006] and Central [Van Erven et al., 2015] stochastic conditions. In both cases the idea behind the assumption is to control the variance of the excess loss of the actions in the neighborhood of the best action.

We do not require that the losses are i.i.d., nor that the Bayes act is in the model. For the Hedge setting it suffices if there is a fixed expert k^* that is always best, i.e. $\mathbb{E}\left[\ell_t^{k^*} \middle| \mathcal{G}_{t-1}\right] = \inf_k \mathbb{E}\left[\ell_t^k \middle| \mathcal{G}_{t-1}\right]$ almost surely for all t. (Here we denote by $\mathcal{G}_{t-1}$ the sigma algebra generated by $\ell_1, \ldots, \ell_{t-1}$, and the *almost surely* quantification refers to the distribution of $\ell_1, \ldots, \ell_{t-1}$.) Similarly, for OCO we assume there is a fixed point $u^* \in \mathcal{U}$ attaining $\min_{u \in \mathcal{U}} \mathbb{E}\left[\ell_t(u) | \mathcal{G}_{t-1}\right]$ at every round t. In either case there may be multiple candidate k^* or u^*. In the succeeding we assume that one is selected. Note that for i.i.d. losses the existence of a minimiser is not such a strong assumption (it is even automatic in the OCO case due to compactness of $\mathcal{U}$), while it is very strong beyond i.i.d. Yet it is not impossible (and actually interesting) as we will show by example in Section 3.

Based on this loss minimiser, we define the *excess losses*, a family of random variables indexed by time $t \in \mathbb{N}$ and expert/point $k \in \mathbb{N}/u \in \mathcal{U}$ as follows

$$x_t^k \;\coloneqq\; \ell_t^k - \ell_t^{k^*} \quad \text{(Hedge)} \qquad \text{and} \qquad x_t^u \;\coloneqq\; \langle u - u^*, \nabla \ell_t(u) \rangle \quad \text{(OCO)}. \tag{3}$$

Note that for the Hedge setting we work with the loss directly. For OCO instead we talk about the linear upper bound on the loss, for this is the quantity that needs to be controlled to make use of the MetaGrad bound (2). With these variables in place, from this point on the story is the same for Hedge and for OCO. So let us write $\mathcal{F}$ for either the set $\mathbb{N}$ of experts or the set $\mathcal{U}$ of points, and f^* for k^* resp. u^*, and let us consider the family $\{x_t^f \mid f \in \mathcal{F}, t \in \mathbb{N}\}$. We call $f \in \mathcal{F}$ *predictors*. The point of these stochastic conditions is that they imply that the variance in the excess loss gets smaller the closer a predictor gets to the optimum in terms of expected excess loss. This is most directly seen in the Bernstein condition:

Condition 1. Fix $B \geq 0$ and $\kappa \in [0,1]$. The family (3) satisfies the (B,κ)-*Bernstein condition* if

$$\mathbb{E}\Big[(x_t^f)^2 \Big| \mathcal{G}_{t-1}\Big] \;\leq\; B\,\mathbb{E}\Big[x_t^f \Big| \mathcal{G}_{t-1}\Big]^{\kappa} \quad \text{almost surely for all } f \in \mathcal{F} \text{ and rounds } t \in \mathbb{N}.$$

Some authors refer to the $\kappa = 1$ case as the *Massart condition*. As shown by Van Erven et al. [2015, Theorem 5.4], for bounded excess losses (which we assume throughout), the Bernstein condition is equivalent to the following condition:

Condition 2. Fix a function $\epsilon : \mathbb{R}_+ \to \mathbb{R}_+$. The family (3) satisfies the ϵ-*central condition* if

$$\frac{1}{\eta} \ln \mathbb{E}\Big[e^{-\eta x_t^f} \Big| \mathcal{G}_{t-1}\Big] \;\leq\; \epsilon(\eta) \qquad \text{almost surely for all } f \in \mathcal{F},\ \eta \geq 0 \text{ and } t \in \mathbb{N}.$$

Van Erven et al. explicitly convert back and forth between the parameters of the Bernstein and Central Condition. In the remainder we will use that Bernstein implies Central (with $\epsilon(\eta) = O((B\eta)^{\frac{1}{1-\kappa}})$). For completeness we include a proof in Appendix B.

2.4 Main Result

In the stochastic case we evaluate the performance of algorithms by $R_T^{f^*}$, i.e. the regret compared to the predictor f^* with minimal expected loss. The expectation $\mathbb{E}[R_T^{f^*}]$ is sometimes called the *pseudo-regret*. The following result shows that second-order methods automatically adapt to the Bernstein condition. (Proof sketch in Section 4.)

Theorem 3. *In any stochastic setting satisfying the (B,κ)-Bernstein Condition 1, the guarantees* (1) *for Squint and* (2) *for MetaGrad imply fast rates for the respective algorithms both in expectation and with high probability. That is,*

$$\mathbb{E}[R_T^{f^*}] \;=\; O\left(K_T^{\frac{1}{2-\kappa}} T^{\frac{1-\kappa}{2-\kappa}}\right),$$

and for any $\delta > 0$, with probability at least $1-\delta$,

$$R_T^{f^*} \;=\; O\left((K_T - \ln\delta)^{\frac{1}{2-\kappa}} T^{\frac{1-\kappa}{2-\kappa}}\right),$$

where for Squint $K_T \coloneqq K_T^{f^*}$ *from* (1) *and for MetaGrad* K_T *is as in* (2).

We see that Squint and MetaGrad (and any other second-order methods achieving the same bounds, as our results only use these bounds and do not depend on the details of the algorithms) adapt automatically to the Bernstein parameters of the distribution, without any tuning. Appendix F provides an extension of Theorem 3 that allows using Squint with uncountable $\mathcal{F}$.

Crucially, the bound provided by Theorem 3 is natural, and, in general, the best one can expect. This can be seen from considering the *statistical learning setting*, which is a special case of our setup. Here (x_t, y_t) are i.i.d. $\sim \mathbb{P}$ and $\mathcal{F}$ is a set of functions from $\mathcal{X}$ to a set of predictions $\mathcal{A}$, with $\ell_t^f \coloneqq \ell(y_t, f(x_t))$ for some loss function $\ell : \mathcal{Y} \times \mathcal{A} \to [0,1]$ such as squared, 0/1, or absolute loss. In this setting one usually considers excess risk, which is the expected loss difference between the learned $\hat{f}$ and the optimal f^*. The minimax expected (over training sample (x^t, y^t)) risk relative to f^* is of order $T^{-1/2}$ (see e.g. Massart and Nédélec [2006], Audibert [2009]). To get better risk rates, one has to impose further assumptions on $\mathbb{P}$. A standard assumption made in such cases is a Bernstein condition with exponent $\kappa > 0$; see e.g. Koltchinskii [2006], Bartlett and Mendelson [2006], or Audibert [2004] or Audibert [2009]; see Van Erven et al. [2015] for how it generalizes the Tsybakov margin and other conditions.

If $\mathcal{F}$ is sufficiently 'simple', e.g. a class with logarithmic entropy numbers (see Appendix F), or, in classification, a VC class, then, if a κ-Bernstein condition holds, ERM (empirical risk minimization) achieves, in expectation, a better excess risk bound of order $O\left((\log T) \cdot T^{-\frac{1}{2-\kappa}}\right)$. The bound interpolates between $T^{-1/2}$ for $\kappa = 0$ and T^{-1} for $\kappa = 1$ (Massart condition). Results of Tsybakov [2004], Massart and Nédélec [2006], Audibert [2009] suggest that this rate can, in general, not be improved upon, and exactly this rate is achieved by ERM and various other algorithms in various settings by e.g. Tsybakov [2004], Audibert [2004, 2009], Bartlett et al. [2006]. By summing from $t = 1$ to T and using ERM at each t to classify the next data point (so that ERM becomes FTL, follow-the-leader), this suggests that we can achieve a cumulative expected regret $\mathbb{E}[R_T^{f^*}]$ of order $O\left((\log T) \cdot T^{\frac{1-\kappa}{2-\kappa}}\right)$. Theorem 3 shows that this is, indeed, also the rate that Squint attains in such cases if $\mathcal{F}$ is countable and the optimal f^* has positive prior mass π^{f^*} (more on this condition below)— we thus see that Squint obtains exactly the rates one would expect from a statistical learning/classification perspective, and the minimax excess risk results in that setting suggests that these cumulative regret rates cannot be improved in general. It was shown earlier by Audibert [2004] that, when equipped with an oracle to tune the learning rate η as a function of t, the rates $O\left((\log T) \cdot T^{\frac{1-\kappa}{2-\kappa}}\right)$ can also be achieved by Hedge, but the exact tuning depends on the unknown κ. Grünwald [2012] provides a means to tune η automatically in terms of the data, but his method — like ERM and all algorithms in the references above — may achieve *linear* regret in worst-case settings, whereas Squint keeps the $O(\sqrt{T})$ guarantee for such cases.

Theorem 3 only gives the desired rate for Squint if $\mathcal{F}$ is countable and $\pi^{f^*} > 0$. The combination of these two assumptions is strong or at least unnatural, and OCO cannot be readily used in all such cases either, so in Appendix F we therefore show how to extend Theorem 3 to the case of infinite $\mathcal{F}$, which can be continuous and thus have $\pi^{f^*} = 0$, as long as $\mathcal{F}$ admits sufficiently small entropy numbers. Incidentally, this also allows us to show that Squint achieves regret rate $O\left((\log T) \cdot T^{\frac{1-\kappa}{2-\kappa}}\right)$ when $\mathcal{F} = \bigcup_{i=1,2,\ldots} \mathcal{F}_i$ is a countably infinite union of $\mathcal{F}_i$ with appropriate entropy numbers; in such cases there can be, at every sample size, a classifier $\hat{f} \in \mathcal{F}$ with 0 empirical error, so that ERM/FTL will always overfit and cannot be used even if the Bernstein condition holds; Squint allows for aggregation of such models. In the remainder of the main text, we concentrate on applications for which Theorem 3 can be used directly, without extensions.

3 Examples

We give examples motivating and illustrating the Bernstein/Central condition for the Hedge and OCO settings. Our examples in the Hedge setting will illustrate Bernstein with $\kappa < 1$ and non i.i.d. distributions. Our OCO examples were chosen to be natural and illustrate fast rates without curvature.

3.1 Hedge Setting: Gap implies Bernstein with $\kappa = 1$

In the Hedge setting, we say that a distribution $\mathbb{P}$ (not necessarily i.i.d.) of expert losses $\{\ell_t^k \mid t, k \in \mathbb{N}\}$ has *gap* $\alpha > 0$ if there is an expert k^* such that

$$\mathbb{E}\left[\ell_t^{k^*} \middle| \mathcal{G}_{t-1}\right] + \alpha \;\leq\; \inf_{k \neq k^*} \mathbb{E}\left[\ell_t^k \middle| \mathcal{G}_{t-1}\right] \qquad \text{almost surely for each round} \in \mathbb{N}.$$

It is clear that the condition can only hold for k^* the minimiser of the expected loss.

Lemma 4. *A distribution with gap α is $(\frac{1}{\alpha}, 1)$-Bernstein.*

Proof. For all k, t we have $\mathbb{E}\left[(x_t^k)^2 \middle| \mathcal{G}_{t-1}\right] \leq 1 = \frac{1}{\alpha}\alpha \leq \frac{1}{\alpha}\mathbb{E}\left[x_t^k \middle| \mathcal{G}_{t-1}\right]$. □

By Theorem 3 we get the $R_T^{k^*} = O(K_T) = O(\ln \ln T)$ rate. Gaillard et al. [2014] show constant regret for finitely many experts and i.i.d. losses with a gap. Our alternative proof above shows that neither finiteness nor i.i.d. are essential for fast rates in this case.

3.2 Hedge Setting: Any $(1, \kappa)$-Bernstein

The next example illustrates that we can sometimes get the fast rates without a gap. And it also shows that we can get any intermediate rate: we construct an example satisfying the Bernstein condition for any $\kappa \in [0, 1]$ of our choosing

(such examples occur naturally in classiciation settings such as those consider in the example in Appendix F).

Fix $\kappa \in [0,1]$. Each expert k is parametrised by a real number $\delta_k \in [0,1/2]$. The only assumption we make is that $\delta_k = 0$ for some k, and $\inf_k\{\delta_k \mid \delta_k > 0\} = 0$. For a concrete example let us choose $\delta_1 = 0$ and $\delta_k = 1/k$. Expert δ has loss $1/2 \pm \delta$ with probability $\frac{1 \pm \delta^{2/\kappa - 1}}{2}$ independently between experts and rounds. Expert δ has mean loss $\frac{1}{2} + \delta^{2/\kappa}$, and so $\delta = 0$ is best, with loss deterministically equal to $1/2$. The squared excess loss of δ is δ^2. So we have the Bernstein condition with exponent κ (but no $\kappa' > \kappa$) and constant 1, and the associated regret rate by Theorem 3.

Note that for $\kappa = 0$ (the hard case) all experts have mean loss equal to $\frac{1}{2}$. So no matter which k^* we designate as the best expert our expected regret is zero. Yet the experts do not agree, as their losses deviate from $\frac{1}{2}$ independently at random. Hence, by the central limit theorem, with high probability our regret is of order $\sqrt{T}$. On the other side of the spectrum, for $\kappa = 1$ (the best case), we do not find a gap. We still have experts arbitrary close to the best expert in mean, but their expected excess loss squared equals their expected excess loss.

ERM/FTL may fail miserably on this type of examples. The clearest case is when $\{k \mid \delta_k > \epsilon\}$ is infinite for some $\epsilon > 0$. Then at any t there will be experts that, by chance, incurred their lower loss every round. Picking any of them will result in expected instantaneous regret at least $\epsilon^{2/\kappa}$, leading to linear regret overall.

The requirement $\delta_k = 0$ for some k is essential. If instead $\delta_k > 0$ for all k then there is no best expert in the class. Theorem 13 in Appendix F shows how to deal with this case.

3.3 Hedge Setting: Markov Chains

Suppose we model a binary sequence $z_1, z_2, \ldots, z_T$ with m-th order Markov chains. As experts we consider all possible functions $f\colon \{0,1\}^m \to \{0,1\}$ that map a history of length m to a prediction for the next outcome, and the loss of expert f is the 0/1-loss: $\ell_t^f = |f(z_{t-m}, \ldots, z_{t-1}) - z_t|$. (We initialize $z_{1-m} = \ldots = z_0 = 0$.) A uniform prior on this finite set of 2^{2^m} experts results in worst-case regret of order $\sqrt{T 2^m}$. Then, if the data are actually generated by an m-th order Markov chain with transition probabilities $\mathbb{P}(z_t = 1 \mid (z_{t-m}, \ldots, z_{t-1}) = \boldsymbol{a}) = p_{\boldsymbol{a}}$, we have $f^*(\boldsymbol{a}) = \mathbf{1}\{p_{\boldsymbol{a}} \geq \frac{1}{2}\}$ and

$$\mathbb{E}\left[(x_t^f)^2 \middle| (z_{t-m}, \ldots, z_{t-1}) = \boldsymbol{a}\right] = 1, \quad \mathbb{E}\left[x_t^f \middle| (z_{t-m}, \ldots, z_{t-1}) = \boldsymbol{a}\right] = 2|p_{\boldsymbol{a}} - \tfrac{1}{2}|$$

for any f such that $f(\boldsymbol{a}) \neq f^*(\boldsymbol{a})$. So the Bernstein condition holds with $\kappa = 1$ and $B = \frac{1}{2 \min_{\boldsymbol{a}} |p_{\boldsymbol{a}} - \frac{1}{2}|}$.

3.4 OCO: Hinge Loss on the Unit Ball

Let $(x_1, y_1), (x_2, y_2), \ldots$ be classification data, with $y_t \in \{-1, +1\}$, and consider the *hinge loss* $\ell_t(u) = \max\{0, 1 - y_t \langle x_t, u \rangle\}$. Now suppose, for simplicity, that

both x_t and u come from the d-dimensional unit Euclidean ball, such that $\langle x_t, u\rangle \in [-1,+1]$ and the hinge is never active, i.e. $\ell_t(u) = 1 - y_t\langle x_t, u\rangle$. Then, if the data turn out to be i.i.d. observations from a fixed distribution $\mathbb{P}$, the Bernstein condition holds with $\kappa = 1$ (The proof can be found in Appendix D):

Lemma 5 (Unregularized Hinge Loss Example)**.** *Consider the hinge loss setting above, where $|\langle x_t, u\rangle| \le 1$. If the data are i.i.d., then the (B,κ)-Bernstein condition is satisfied with $\kappa = 1$ and $B = \frac{2\lambda_{max}}{\|\boldsymbol{\mu}\|}$, where λ_{max} is the maximum eigenvalue of $\mathbb{E}[\boldsymbol{X}\boldsymbol{X}^\top]$ and $\boldsymbol{\mu} = \mathbb{E}[Y\boldsymbol{X}]$, provided that $\|\boldsymbol{\mu}\| > 0$.*

In particular, if $\boldsymbol{X}_t$ is uniformly distributed on the sphere and $Y_t = \operatorname{sign}(\langle \bar{\boldsymbol{u}}, \boldsymbol{X}_t\rangle)$ is the noiseless classification of $\boldsymbol{X}_t$ according to the hyperplane with normal vector $\bar{\boldsymbol{u}}$, then $B \le \frac{c}{\sqrt{d}}$ for some absolute constant $c > 0$.

The excluded case $\|\boldsymbol{\mu}\| = 0$ only happens in the degenerate case that there is nothing to learn, because $\boldsymbol{\mu} = \boldsymbol{0}$ implies that the expected hinge loss is 1, its maximal value, for all $\boldsymbol{u}$.

3.5 OCO: Absolute Loss

Let $\mathcal{U} = [0,1]$ be the unit interval. Consider $\ell_t(u) = |u - x_t|$ where $x_t \in [0,1]$ are drawn i.i.d. from $\mathbb{P}$. Let $u^* \in \arg\min_u \mathbb{E}|u - x|$ minimize the expected loss. In this case we may simplify $\langle w - u^*, \nabla\ell(w)\rangle = (w - u^*)\operatorname{sign}(w - x)$. To satisfy the Bernstein condition, we therefore want B such that, for all $w \in [0,1]$,

$$\mathbb{E}\Big[\big((w-u^*)\operatorname{sign}(w-x)\big)^2\Big] \;\le\; B\,\mathbb{E}\big[(w-u^*)\operatorname{sign}(w-x)\big]^\kappa .$$

That is,

$$|w-u^*|^{2-\kappa} \;\le\; B2^\kappa\big|\mathbb{P}(x \le w) - \tfrac{1}{2}\big|^\kappa .$$

For instance, if the distribution of x has a strictly positive density $p(x) \ge m > 0$, then u^* is the median and $|\mathbb{P}(x \le w) - \frac{1}{2}| = |\mathbb{P}(x \le w) - \mathbb{P}(x \le u^*)| \ge m|w - u^*|$, so the condition holds with $\kappa = 1$ and $B = \frac{1}{2m}$. Alternatively, for a discrete distribution on two points a and b with probabilities p and $1-p$, the condition holds with $\kappa = 1$ and $B = \frac{1}{|2p-1|}$, provided that $p \ne \frac{1}{2}$, as can be seen by bounding $|w - u^*| \le 1$ and $|\mathbb{P}(x \le w) - \frac{1}{2}| \ge |p - \frac{1}{2}|$.

4 Proof of Main Result

This section builds up to prove our main result Theorem 3. We first introduce a handy abbreviation that allows us to reason simultaneously in expectation and with high probability. We then identify the minimal ϵ for which the Central Condition 2 holds. We then show how we can introduce a second-order adjustment, and characterize the cost. Combination with either worst-case regret bound then yields the desired result.

4.1 Notation: Exponential Stochastic Negativity and Inequality

We introduce a convenient shorthand notation that we will use throughout this paper.

Definition 6. A random variable X is *exponentially stochastically negative*, denoted $X \trianglelefteq 0$, if $\mathbb{E}[e^X] \leq 1$. For any $\eta \geq 0$, we write $X \trianglelefteq_\eta 0$ if $\eta X \trianglelefteq 0$. For any pair of random variables X and Y, we say that X is *exponentially stochastically less* than Y, denoted $X \trianglelefteq Y$, if $X - Y \trianglelefteq 0$.

Lemma 7. *Exponential stochastic negativity has the following useful properties:*

1. *(Negativity). Let $X \trianglelefteq 0$. As the notation suggests X is negative in expectation and with high probability. That is $\mathbb{E}[X] \leq 0$ and $\mathbb{P}\{X \geq -\ln\delta\} \leq \delta$ for all $\delta > 0$.*
2. *(Convex combination). Let $\left\{X^f\right\}_{f\in\mathcal{F}}$ be a family of random variables and let w be a distribution on $\mathcal{F}$. If $X^f \trianglelefteq 0$ for all f then $\mathbb{E}_{f\sim w}[X^f] \trianglelefteq 0$.*
3. *(Chain rule). Let $X_1, X_2, \ldots$ be adapted to filtration $\mathcal{G}_1 \subseteq \mathcal{G}_2 \ldots$ (i.e. X_t is $\mathcal{G}_t$-measurable for each t). If $X_t|\mathcal{G}_{t-1} \trianglelefteq 0$ almost surely for all t, then $\sum_{t=1}^T X_t \trianglelefteq 0$ for all $T \geq 0$.*

Proof. Negativity: By Jensen's inequality $\mathbb{E}[X] \leq \ln\mathbb{E}\left[e^X\right] \leq 0$, whereas by Markov's inequality $\mathbb{P}\{X \geq -\ln\delta\} = \mathbb{P}\left\{e^X \geq \frac{1}{\delta}\right\} \leq \delta\,\mathbb{E}\left[e^X\right] \leq \delta$. *Convex combination*: By Jensen's inequality $\mathbb{E}\left[e^{\mathbb{E}_{f\sim w}[X^f]}\right] \leq \mathbb{E}_{f\sim w}\mathbb{E}\left[e^{X^f}\right] \leq 1$. *Chain rule*: By induction. The base case $T = 0$ holds trivially. For $T > 0$ we have $\mathbb{E}\left[e^{\sum_{t=1}^T X_t}\right] = \mathbb{E}\left[e^{\sum_{t=1}^{T-1} X_t}\,\mathbb{E}\left[e^{X_T}\middle|\mathcal{G}_{T-1}\right]\right] \leq \mathbb{E}\left[e^{\sum_{t=1}^{T-1} X_t}\right] \leq 1$. □

4.2 Normalized Cumulant Generating Function

To prove our main result we will make use of the Central Condition 2. For any distribution $\mathbb{P}$ this condition will hold for some ϵ (which may be trivial). In this section we construct the smallest ϵ for which it holds and derive a few useful properties of that ϵ.

Consider the family (3) of excess loss variables x_t^f. We assume that $x_t^f \in [-1, 1]$ is bounded in a range of width 2 and has positive mean $\mathbb{E}[x_t^f|\mathcal{G}_{t-1}] \geq 0$ by definition of f^*. As we will see, the complexity of our learning problem will be governed by the distribution of x_t^f. In particular, we will look at the *normalized cumulant generating function* for $\eta \geq 0$:

$$\epsilon_t^f(\eta) \;\coloneqq\; \frac{1}{\eta}\ln\mathbb{E}\left[e^{-\eta x_t^f}\middle|\mathcal{G}_{t-1}\right]$$

By construction $-x_t^f \trianglelefteq_\eta \epsilon_t^f(\eta)$. Boundedness of $x_t^f \in [-1, 1]$ immediately results in $\epsilon_t^f(\eta) \in [-1, 1]$. Moreover, Hoeffding's inequality (see e.g. Cesa-Bianchi and Lugosi

[2006, Lemma 2.2]) tells us that $\epsilon_t^f(\eta) \le \eta/2$ while Jensen's inequality gives $\epsilon_t^f(\eta) \ge -\mathbb{E}\left[x_t^f \middle| \mathcal{G}_{t-1}\right]$. The dual representation $\epsilon_t^f(\eta) = \sup_Q -\mathbb{E}_Q[x] - \frac{1}{\eta}\mathrm{KL}\big(Q(x) \big\| \mathbb{P}(x_t^f|\mathcal{G}_{t-1})\big)$ reveals that $\epsilon_t^f(\eta)$ is increasing in η. The value at $\eta = 0$ is obtained by continuity from $\epsilon_t^f(0) \coloneqq \lim_{\eta\to 0}\epsilon_t^f(\eta) = -\mathbb{E}[x_t^f|\mathcal{G}_{t-1}] \le 0$.

To get a uniform control over the class $\mathcal{F}$, we will make use of the maximum

$$\epsilon_t(\eta) \;\coloneqq\; \sup_{f\in\mathcal{F}} \epsilon_t^f(\eta) \qquad \text{and} \qquad \epsilon(\eta) \;\coloneqq\; \sup_t\; \epsilon_t(\eta). \tag{4}$$

The functions ϵ_t and ϵ inherit most properties of each ϵ_t^f, but in addition since $f^* \in \mathcal{F}$ and $\epsilon_t^{f^*}(\eta) = 0$, we see that $\epsilon_t(\eta) \ge 0$ and also that $\epsilon_t(0) = 0$. Moreover, since $\epsilon_t(\eta) \le \eta/2$ we have $\lim_{\eta\to 0}\epsilon_t(\eta) = 0$. In this paper we will judge the complexity of the interplay of the generating distribution $\mathbb{P}$ with the class $\mathcal{F}$ by how $\epsilon_t(\eta) \to 0$ as $\eta \to 0$. By construction the Central Condition 2 holds with $\epsilon(\eta)$.

4.3 A Second-order Adjustment to Exponential Stochastic Inequality

We now show a technical lemma showing that, roughly, the square of a bounded Central random variable is exponentially stochastically less than that variable itself. Consider any random variable $x \in [-1, 1]$, and let us denote its normalized cumulant generating function by $\epsilon(\eta) = \frac{1}{\eta}\ln\mathbb{E}[e^{-\eta x}]$. (In particular, see Definition 6, $-x \trianglelefteq_\eta \epsilon(\eta)$ for all $\eta \ge 0$.) Intuitively, small $\epsilon(\eta) \ll \eta/2$ is special, indicating that x cannot be often very negative. The following lemma shows that if x is special to degree $\epsilon(\eta)$, then the smaller quantity $\approx x - \frac{\eta}{2}x^2$ is also special at only mildly weaker degree $\approx \epsilon(2\eta)$.

Lemma 8. *For any random variable $x \in [-1, 1]$ and any $\eta \ge 0$*

$$\frac{1}{\eta}\ln\mathbb{E}\left[e^{c\eta^2x^2-\eta x}\right] \;\le\; \epsilon(2\eta) + c\eta\epsilon(2\eta)^2 \qquad \textit{where} \qquad c \;=\; \frac{1}{1+\sqrt{1+4\eta^2}}.$$

In the notation of Section 4.1, the lemma reads

$$-x \;\trianglelefteq_\eta\; \epsilon(\eta) \quad \text{for all } \eta \ge 0 \quad \text{implies} \quad c\eta x^2 - x \;\trianglelefteq_\eta\; \epsilon(2\eta) + c\eta\epsilon(2\eta)^2 \quad \text{for all } \eta \ge 0.$$

Proof. By Theorem 12 in Appendix A with $\gamma = 2\eta$ and the largest admissible c for (5). □

Note that for $\eta = 0$ the lemma trivializes, telling us $-\mathbb{E}[x] \le \epsilon(0)$ where we have in fact equality. Note also that the right-hand side is an increasing function in $\epsilon(2\eta)$ (the quadratic in $\epsilon(2\eta)$ has positive derivative for all $\epsilon(2\eta) \ge -\frac{1+\sqrt{4\eta^2+1}}{2\eta} < -1$.)

4.4 From Second-order Bound to Bound in Terms of Parameter of Distribution

The next step toward fast rates is to obtain from a second-order bound, which involves the algorithm, another bound strictly in terms of the parameters of the distribution, which do not refer to the algorithm. The proof is in Appendix C.

Theorem 9. *Consider either Squint in the Hedge setting or MetaGrad for OCO. Let* $\left\{x_t^f \middle| f \in \mathcal{F}\right\}$ *be the associated the excess loss family from* (3)*, and let* $\epsilon(\eta)$ *be, as in* (4)*, the corresponding maximal normalized cumulant generating function. For the Hedge setting let* $K_T \coloneqq K_T^{f^*}$ *as in* (1)*, for OCO let* K_T *be as in* (2)*. Then for each* $\gamma \geq 0$ *with* c *as in Lemma 8,*

$$R_T^{f^*} \trianglelefteq_\gamma \frac{K_T}{c\gamma} + T\epsilon(2\gamma)(1 + c\gamma^2) + 2K_T.$$

To prove our main Theorem 3 we invoke the Bernstein Condition 1 to bound $\epsilon(2\epsilon)$ as a polynomial in γ, and then tune γ to optimize the bound. The details of the proof can be found in Appendix E.

5 Conclusion

We show that it is possible for online learning methods to provide both the safety and robustness of a worst-case regret bound and be adaptive to favorable stochastic environments. We focus on Squint and MetaGrad, methods for online learning with individual sequence regret guarantees of a particular second order form. We show that such guarantees imply automatic adaptivity to the Bernstein parameters of stochastic environments, and result in the corresponding fast regret rates.

References

Jean-Yves Audibert. *PAC-Bayesian statistical learning theory*. PhD thesis, Université Paris VI, 2004.

Jean-Yves Audibert. Fast learning rates in statistical inference through aggregation. *The Annals of Statistics*, 37(4):1591–1646, 2009.

Peter L. Bartlett and Shahar Mendelson. Empirical minimization. *Probability Theory and Related Fields*, 135(3):311–334, 2006.

Peter L Bartlett, Michael I Jordan, and Jon D McAuliffe. Convexity, classification, and risk bounds. *Journal of the American Statistical Association*, 101 (473):138–156, 2006.

Nicolò Cesa-Bianchi and Gábor Lugosi. *Prediction, learning, and games*. Cambridge University Press, 2006.

Nicolò Cesa-Bianchi, Yishay Mansour, and Gilles Stoltz. Improved second-order bounds for prediction with expert advice. *Machine Learning*, 66(2/3):321–352, 2007.

Chao-Kai Chiang, Tianbao Yang, Chia-Jung Le, Mehrdad Mahdavi, Chi-Jen Lu, Rong Jin, and Shenghuo Zhu. Online optimization with gradual variations. In *Proceedings of the 25th Conference on Learning Theory (COLT)*, pages 6.1–6.20, 2012.

Koby Crammer, Alex Kulesza, and Mark Dredze. Adaptive regularization of weight vectors. In *Advances in Neural Information Processing Systems 22*, pages 414–422, 2009.

John Duchi, Elad Hazan, and Yoram Singer. Adaptive subgradient methods for online learning and stochastic optimization. *Journal of Machine Learning Research*, 12:2121–2159, 2011.

Tim van Erven, Peter D. Grünwald, Nishant A. Mehta, Mark D. Reid, and Robert C. Williamson. Fast rates in statistical and online learning. *Journal of Machine Learning Research*, 16:1793–1861, 2015.

Yoav Freund and Robert E. Schapire. A decision-theoretic generalization of on-line learning and an application to boosting. *Journal of Computer and System Sciences*, 55:119–139, 1997.

Pierre Gaillard and Sébastien Gerchinovitz. A chaining algorithm for online nonparametric regression. In *Proceedings of the 28th Conference on Learning Theory (COLT)*, 2015.

Pierre Gaillard, Gilles Stoltz, and Tim van Erven. A second-order bound with excess losses. In *JMLR Workshop and Conference Proceedings*, volume 35: Proceedings of the 27th Conference on Learning Theory (COLT), pages 176–196, 2014.

Peter Grünwald. The safe Bayesian: learning the learning rate via the mixability gap. In *Proceedings 23rd International Conference on Algorithmic Learning Theory (ALT '12)*. Springer, 2012.

Elad Hazan and Satyen Kale. Extracting certainty from uncertainty: Regret bounded by variation in costs. *Machine learning*, 80(2-3):165–188, 2010.

Vladimir Koltchinskii. Local Rademacher complexities and oracle inequalities in risk minimization. *The Annals of Statistics*, 34(6):2593–2656, 2006. Based on IMS (Institute of Mathematical Statistics) Medaillion Lecture.

Wouter M. Koolen. The relative entropy bound for Squint. Blog entry on `http://blog.wouterkoolen.info/`, August 2015.

Wouter M. Koolen and Tim van Erven. Second-order quantile methods for experts and combinatorial games. In *Proceedings of the 28th Conference on Learning Theory (COLT)*, pages 1155–1175, 2015.

Wouter M. Koolen and Tim van Erven. Metagrad: Faster convergence without curvature in online convex optimization. *ArXiv:1604.08740*, April 2016.

Wouter M. Koolen, Tim van Erven, and Peter D. Grünwald. Learning the learning rate for prediction with expert advice. In *Advances in Neural Information Processing Systems 27*, pages 2294–2302. 2014.

Haipeng Luo and Robert E. Schapire. Achieving all with no parameters: Adaptive normalhedge. In *Proceedings of the 28th Conference on Learning Theory (COLT)*, pages 1286–1304, 2015.

Pascal Massart and Élodie Nédélec. Risk bounds for statistical learning. *The Annals of Statistics*, 34(5):2326–2366, 2006.

H. Brendan McMahan and Matthew J. Streeter. Adaptive bound optimization for online convex optimization. In *Proceedings of the 23rd Conference on Learning Theory (COLT)*, pages 244–256, 2010.

Nishant A Mehta and Robert C Williamson. From stochastic mixability to fast rates. In *Advances in Neural Information Processing Systems 27*, pages 1197–1205. Curran Associates, Inc., 2014.

Alexander Rakhlin and Karthik Sridharan. Online nonparametric regression. In *Proceedings of the 27th Conference on Learning Theory (COLT)*, 2014.

Steven de Rooij, Tim van Erven, Peter D. Grünwald, and Wouter M. Koolen. Follow the leader if you can, Hedge if you must. *Journal of Machine Learning Research*, 15:1281–1316, April 2014.

Amir Sani, Gergely Neu, and Alessandro Lazaric. Exploiting easy data in online optimization. In *Advances in Neural Information Processing Systems 27*, pages 810–818. 2014.

Shai Shalev-Shwartz. Online learning and online convex optimization. *Foundations and Trends in Machine Learning*, 4(2):107–194, 2011.

Jacob Steinhardt and Percy Liang. Adaptivity and optimism: An improved exponentiated gradient algorithm. In *Proceedings of the 31th Annual International Conference on Machine Learning (ICML)*, pages 1593–1601, 2014.

Alexandre B. Tsybakov. Optimal aggregation of classifiers in statistical learning. *Annals of Statistics*, 32:135–166, 2004.

Olivier Wintenberger. Optimal learning with Bernstein Online Aggregation. *ArXiv:1404.1356*, 2015.

A Second-order Adjustment of Exponential Stochastic Inequality

In this section we prove a stronger form of Lemma 8. We would like to remark that our solution to this problem was inspired by the general moments problem studied by Mehta and Williamson [2014, Section 3], especially because this connection became invisible during the simplification of our proofs.

We will be thinking about two learning rates, $0 \le \eta \le \gamma$. The larger one, γ, will be where we evaluate $\epsilon(\gamma)$. So γ controls the strength of the assumption. The smaller one, η, will be the learning rate at which we obtain the conclusion. The point is to get a large amount of quadratic x^2 in the conclusion, as governed by the constant c. Obviously, the more greedy we are in η and γ, the smaller the c for which we can get any traction. This trade-off is captured by the following relationship between γ, η and c that we will make use of throughout this section.

$$0 \;\le\; c \;\le\; \frac{\sqrt{2|2\eta-\gamma|+\gamma^2+1}-|2\eta-\gamma|-1}{4\eta^2} \tag{5}$$

Positivity of c is not that important, as the desired inequality is trivial for $c \le 0$. The following inequality is useful later.

Lemma 10. *Let $0 \le \eta \le \gamma$ and c satisfy (5). Then*

$$1 \;\ge\; 2c\eta$$

Proof. We need to show

$$2\eta \;\ge\; \sqrt{2|2\eta-\gamma|+\gamma^2+1}-|2\eta-\gamma|-1$$

that is

$$(2\eta+|2\eta-\gamma|+1)^2 \;\ge\; 2|2\eta-\gamma|+\gamma^2+1$$

Expanding the left-hand side square results in

$$4\eta^2+4\eta|2\eta-\gamma|+4\eta+|2\eta-\gamma|^2+2|2\eta-\gamma|+1 \;=\; 4\eta(2\eta-\gamma)+4\eta|2\eta-\gamma|+4\eta+\gamma^2+2|2\eta-\gamma|+1$$

which definitely exceeds the right-hand side above. □

We now put our assumption to use. In the following Lemma we show that it implies a not-in-expectation-but-with-a-correction-term version of the result we are after.

Lemma 11. *Fix $0 \le \eta \le \gamma$ and let c satisfy (5). Then for each $x \in [-1,1]$ and $\epsilon \in [-1,1]$ we have*

$$e^{c\eta^2x^2-\eta x}-\frac{e^{-\gamma(x+\epsilon)}-1}{\gamma}\eta(1+2c\eta\epsilon)e^{c\eta^2\epsilon^2+\eta\epsilon} \;\le\; e^{c\eta^2\epsilon^2+\eta\epsilon}.$$

Proof. We will show that the left-hand side is maximized over $x \in [-1,1]$ at $x = -\epsilon$. First, its derivative equals

$$e^{-\gamma x}\eta\big(h(-\epsilon) - h(x)\big) \qquad \text{where} \qquad h(x) = (1 - 2c\eta x)e^{c\eta^2 x^2 + (\gamma-\eta)x}.$$

This indeed equals zero at $x = -\epsilon$. To show that $x = -\epsilon$ is indeed a maximum and that there are no other maxima it suffices to show that $h(x)$ is increasing on $x \in [-1,1]$. We have

$$h'(x) \;=\; \left(-4c^2\eta^3 x^2 + 2c\eta x(2\eta - \gamma) - 2c\eta + \gamma - \eta\right) e^{c\eta^2 x^2 + (\gamma-\eta)x}$$

As the term in parentheses is concave in x, it suffices to show that $h'(x) \geq 0$ for $x \in \{-1,1\}$, i.e.

$$-4c^2\eta^3 - 2c\eta|2\eta - \gamma| - 2c\eta + \gamma - \eta \;\geq\; 0$$

Solving the quadratic in c, we see that this holds if (5), as required. □

Finally, we are ready for the general version of the claim.

Theorem 12. *Pick $0 \leq \eta \leq \gamma$ and c satisfying (5). Let $\epsilon \in [-1,1]$. Then for any $x \in [-1,1]$ with $\mathbb{E}\, e^{-\gamma x} \leq e^{\gamma\epsilon}$ we have*

$$\mathbb{E}\, e^{c\eta^2 x^2 - \eta x} \;\leq\; e^{c\eta^2\epsilon^2 + \eta\epsilon}.$$

Proof. Taking expectation over Lemma 11, we find

$$\mathbb{E}\, e^{c\eta^2 x^2 - \eta x} \;\leq\; e^{c\eta^2\epsilon^2 + \eta\epsilon} + \frac{\mathbb{E}\, e^{-\gamma(x+\epsilon)} - 1}{\gamma}\eta(1 + 2c\eta\epsilon)e^{c\eta^2\epsilon^2 + \eta\epsilon},$$

and the claim follows by bounding the right-most term by 0. (Note that the factor $1 + 2c\eta\epsilon$ is positive by Lemma 10.) □

B Bernstein to Central

An indexed family of random variables $\{x^f \mid f \in \mathcal{F}\}$ satisfies the (B,κ)-*Bernstein condition* if

$$\mathbb{E}\left[(x^f)^2\right] \;\leq\; B\,\mathbb{E}\left[x^f\right]^\kappa \qquad \text{for all } f \in \mathcal{F}$$

and it satisfies the ϵ-*Central condition* if

$$-x^f \;\trianglelefteq_\eta\; \epsilon(\eta) \qquad \text{for all } f \in \mathcal{F} \text{ and } \eta \geq 0$$

We now show that Bernstein implies Central. This is a special case of [Van Erven et al., 2015, Theorem 5.4, Part 1]. Assume Bernstein. Then by the Bernstein Sandwich [Koolen et al., 2014, Lemma C.2], simplifying $e^\eta - \eta - 1 \leq \eta^2$, which holds for small enough $\eta \leq 1.79328$, and by the Bernstein assumption

$$\epsilon^f(\eta) \;=\; \frac{1}{\eta}\ln\mathbb{E}\left[e^{-\eta x^f}\right] \;\leq\; \eta\mathbb{E}[(x^f)^2] - \mathbb{E}[x^f] \;\leq\; \eta B\,\mathbb{E}[x^f]^\kappa - \mathbb{E}[x^f]$$

Then (the maximizer is found at $x = (B\eta\kappa)^{\frac{1}{1-\kappa}}$ (which is ≤ 1 when $B\kappa\eta \le 1$, so for small enough η this is a reasonable point))

$$\epsilon(\eta) \;=\; \sup_f \epsilon^f(\eta) \;\le\; \sup_x \eta B x^\kappa - x \;=\; \frac{1-\kappa}{\kappa}(B\eta\kappa)^{\frac{1}{1-\kappa}}.$$

C Proof of Theorem 9

Proof. For the Hedge setting, let us write $x_t \coloneqq \mathbb{E}_{k\sim w_t}[x_t^k]$ for the excess loss of the learner in round t. Then $x_t \in [-1,1]$ and $-x_t \trianglelefteq_\eta \epsilon(\eta)$ by Lemma 7. Now by definition of x_t^k in (3) and R_T^k and V_T^k (see Section 2.1)

$$\sum_{t=1}^T x_t \;=\; \sum_{t=1}^T \left(\langle w_t, \ell_t\rangle - \ell_t^{k^*}\right) \;=\; R_T^{k^*} \quad \text{and}$$
$$\sum_{t=1}^T (x_t)^2 \;=\; \sum_{t=1}^T \left(\langle w_t, \ell_t\rangle - \ell_t^{k^*}\right)^2 \;=\; V_T^{k^*}.$$

For OCO, let us write $x_t \coloneqq x_t^{w_t}$ for the excess loss of the learner in round t. Again $x_t \in [-1,1]$ and we have $-x_t \trianglelefteq_\eta \epsilon(\eta)$ by construction of $\epsilon(\eta)$. Moreover, from the definition of $\tilde{R}_T^u$ and V_T^u from Section 2.2,

$$\sum_{t=1}^T x_t \;=\; \sum_{t=1}^T \langle w_t - u^*, \nabla\ell_t(w_t)\rangle \;=\; \tilde{R}_T^{u^*} \quad \text{and}$$
$$\sum_{t=1}^T (x_t)^2 \;=\; \sum_{t=1}^T \langle w_t - u^*, \nabla\ell_t(w_t)\rangle^2 \;=\; V_T^{u^*}.$$

With this notation the second-order regret bounds (1) and (2) both state

$$\sum_{t=1}^T x_t \;\le\; \sqrt{\left(\sum_{t=1}^T x_t^2\right) K_T} + K_T. \tag{6}$$

Now fix $\gamma \ge 0$. For any t, as $-x_t \trianglelefteq_\eta \epsilon(\eta)$, Lemma 8 gives

$$c\gamma x_t^2 - x_t \;\trianglelefteq_\gamma\; \epsilon(2\gamma)(1 + c\gamma^2)$$

By telescoping over rounds (using the chain rule Lemma 7), we obtain

$$c\gamma \sum_{t=1}^T x_t^2 - \sum_{t=1}^T x_t \;\trianglelefteq_\gamma\; T\epsilon(2\gamma)(1 + c\gamma^2). \tag{7}$$

The individual sequence regret bound (6) gives us (since $2\sqrt{ab} = \inf_\eta \eta a + b/\eta$) for every $\eta \ge 0$

$$\sum_{t=1}^T x_t \;\le\; \frac{\eta}{2}\sum_{t=1}^T x_t^2 + \frac{K_T}{2\eta} + K_T.$$

Plugging in $\eta = c\gamma$ (this implies $\eta \in [0, 1/2]$ and $\gamma = \frac{2\eta}{1-4\eta^2}$) and combination with (7) results in

$$\sum_{t=1}^{T} x_t \ \trianglelefteq_\gamma \ \frac{K_T}{c\gamma} + T\epsilon(2\gamma)(1 + c\gamma^2) + 2K_T.$$

For the Hedge setting this proves the theorem. For OCO we finish with $R_T^{u^*} \leq \tilde{R}_T^{u^*}$. □

D Proof of Lemma 5

Proof. Since, by assumption, $\boldsymbol{u}$ and $\boldsymbol{X}$ have length at most 1, the hinge loss simplifies to $\ell(\boldsymbol{u}) = 1 - Y\langle \boldsymbol{u}, \boldsymbol{X}\rangle$ with gradient $\nabla\ell(\boldsymbol{u}) = -Y\boldsymbol{X}$. This implies that

$$\boldsymbol{u}^* \coloneqq \arg\min_{\boldsymbol{u}} \mathbb{E}\left[\ell(\boldsymbol{u})\right] = \frac{\boldsymbol{\mu}}{\|\boldsymbol{\mu}\|}, \tag{8}$$

and

$$\begin{aligned}(\boldsymbol{w} - \boldsymbol{u}^*)^\top \mathbb{E}\left[\nabla\ell(\boldsymbol{w})\nabla\ell(\boldsymbol{w})^\top\right](\boldsymbol{w} - \boldsymbol{u}^*) &= (\boldsymbol{w} - \boldsymbol{u}^*)^\top \mathbb{E}\left[\boldsymbol{X}\boldsymbol{X}^\top\right](\boldsymbol{w} - \boldsymbol{u}^*) \\ &\leq \lambda_{\max}(\boldsymbol{w} - \boldsymbol{u}^*)^\top(\boldsymbol{w} - \boldsymbol{u}^*) \leq 2\lambda_{\max}(1 - \langle \boldsymbol{w}, \boldsymbol{u}^*\rangle) \\ &= \frac{2\lambda_{\max}}{\|\boldsymbol{\mu}\|}(\boldsymbol{w} - \boldsymbol{u}^*)^\top(-\boldsymbol{\mu}) = \frac{2\lambda_{\max}}{\|\boldsymbol{\mu}\|}(\boldsymbol{w} - \boldsymbol{u}^*)^\top \mathbb{E}\left[\nabla\ell(\boldsymbol{w})\right],\end{aligned}$$

which proves the first part of the lemma

For the second part, we first observe that $\lambda_{\max} = 1/d$. Then, to compute $\|\boldsymbol{\mu}\|$, assume without loss of generality that $\|\bar{\boldsymbol{u}}\| = 1$, in which case $\bar{\boldsymbol{u}} = \boldsymbol{u}^*$. Now symmetry of the distribution of $\boldsymbol{X}$ conditional on $\langle \boldsymbol{X}, \boldsymbol{u}^*\rangle$ gives

$$\begin{aligned}&\mathbb{E}\left[Y\boldsymbol{X} \mid \langle \boldsymbol{X}, \boldsymbol{u}^*\rangle\right] \\ &\quad = \operatorname{sign}(\langle \boldsymbol{X}, \boldsymbol{u}^*\rangle)\,\mathbb{E}\left[\boldsymbol{X} \mid \langle \boldsymbol{X}, \boldsymbol{u}^*\rangle\right] = \operatorname{sign}(\langle \boldsymbol{X}, \boldsymbol{u}^*\rangle)\langle \boldsymbol{X}, \boldsymbol{u}^*\rangle\boldsymbol{u}^* = |\langle \boldsymbol{X}, \boldsymbol{u}^*\rangle|\boldsymbol{u}^*.\end{aligned}$$

By rotational symmetry, we may further assume without loss of generality that $\boldsymbol{u}^* = \boldsymbol{e}_1$ is the first unit vector in the standard basis, and therefore

$$\|\boldsymbol{\mu}\| = \|\,\mathbb{E}\left[|\langle \boldsymbol{X}, \boldsymbol{u}^*\rangle|\right]\boldsymbol{u}^*\| = \mathbb{E}\left[|X_1|\right].$$

If $\boldsymbol{Z} = (Z_1, \ldots, Z_d)$ is multivariate Gaussian $\mathcal{N}(0, I)$. Then $\boldsymbol{X} = \boldsymbol{Z}/\|\boldsymbol{Z}\|$ is uniformly distributed on the sphere, so

$$\mathbb{E}[|X_1|] = \mathbb{E}\left[\frac{|Z_1|}{\|\boldsymbol{Z}\|}\right] \geq \frac{1}{4\sqrt{d}}\,\mathbb{P}\left(|Z_1| \geq \tfrac{1}{2} \wedge \|\boldsymbol{Z}\| \leq 2\sqrt{d}\right).$$

Since $\mathbb{P}\left(|Z_1| < \frac{1}{2}\right) \leq 0.4$ and $\mathbb{P}\left(\|\boldsymbol{Z}\| \geq 2\sqrt{d}\right) \leq \frac{1}{4d}\,\mathbb{E}\left[\|\boldsymbol{Z}\|^2\right] = \frac{1}{4}$, we have

$$\mathbb{P}\left(|Z_1| \geq \tfrac{1}{2} \wedge \|\boldsymbol{Z}\| \leq 2\sqrt{d}\right) \geq 1 - 0.4 - \frac{1}{4} = 0.35,$$

from which the conclusion of the second part follows with $c = 8/0.35$. □

E Proof of Theorem 3

Proof. As pointed out below Condition 2, the (B,κ)-Bernstein condition implies the central condition with $\epsilon(\eta) \le (\eta B)^{\frac{1}{1-\kappa}}$. By Theorem 9, using $1/c \le 2(1+\gamma^2)$ and $c \le \frac{1}{2}$, we find that for all $\gamma \ge 0$

$$R_T^{f^*} \trianglelefteq_\gamma (1+\gamma^2)\frac{2K_T}{\gamma} + (1+\frac{1}{2}\gamma^2)\epsilon(2\gamma)T + 2K_T. \tag{9}$$

By Lemma 7 this implies for all $\gamma \ge 0$

$$\mathbb{E}[R_T^{f^*}] \le (1+\gamma^2)\frac{2K_T}{\gamma} + (1+\frac{1}{2}\gamma^2)\epsilon(2\gamma)T + 2K_T.$$

It remains to tune γ to exploit the stochastic condition expressed by ϵ. Reducing the above right-hand-side expression to its main terms and setting the derivative to zero suggests picking

$$\hat{\gamma} = \left(\frac{2K_T(1-\kappa)(2B)^{-\frac{1}{1-\kappa}}}{T}\right)^{\frac{1-\kappa}{2-\kappa}} = O\Big(\big(K_T/T\big)^{\frac{1-\kappa}{2-\kappa}}\Big).$$

Plugging in this tuning, we find

$$\mathbb{E}[R_T^{f^*}] \le (2-\kappa)\,(4K_TB)^{\frac{1}{2-\kappa}}\,(T/(1-\kappa))^{\frac{1-\kappa}{2-\kappa}} + (5-\kappa)K_T$$

Finally, using that $(2-\kappa)\,(4B)^{\frac{1}{2-\kappa}}\,(1/(1-\kappa))^{\frac{1-\kappa}{2-\kappa}}$ is maximized in κ at $\kappa = 1-\frac{1}{4B}$ where it takes value $1+4B$, we may simplify this to

$$\mathbb{E}[R_T^{f^*}] \le (1+4B)\,(K_T/4)^{\frac{1}{2-\kappa}}T^{\frac{1-\kappa}{2-\kappa}} + (5-\kappa)K_T = O\Big(K_T^{\frac{1}{2-\kappa}}T^{\frac{1-\kappa}{2-\kappa}}\Big),$$

which gives the first claim of the Theorem. Lemma 7 applied to (9) also implies that for all $\delta \ge 0$ with probability at least $1-\delta$

$$R_T^{f^*} \le (1+\gamma^2)\frac{2K_T}{\gamma} + (1+\frac{1}{2}\gamma^2)\epsilon(2\gamma)T + 2K_T + \frac{-\ln\delta}{\gamma}.$$

Tuning γ as before with K_T replaced by $K_T + \frac{-\ln\delta}{2}$ yields the second claim. □

F Continuous Models

We now consider Squint with models of predictors $\mathcal{F}$ that have uncountably many elements so that in general $\pi^{f^*} = 0$, and each $f \in \mathcal{F}$ is a function from $\mathcal{X}$ to $\mathcal{A}$, with $\ell_t^f := \ell(y_t, f(x_t))$ for some fixed loss function $\ell : \mathcal{Y}\times\mathcal{A} \to [0,1]$. This setting includes standard parametric models in classification and regression but also countable unions thereof as well as nonparametric models. We first present an extension of Theorem 3 to this case; we then give an illustration of this result with sup-norm metric entropy numbers.

Squint can be straightforwardly applied to uncountable models, but now the weight vector w_t output by Squint at time t takes the form of a distribution on the set $\mathcal{F}$. For general distributions u on $\mathcal{F}$, the loss u incurs at time t is now defined as $\ell_t^u := \mathbb{E}_{f\sim u}[\ell_t^f]$, so that the loss of Squint at time t is given by $\ell_t^{w_t}$, which generalizes the expression $\langle w_t, \ell_t\rangle$ for the countable case. The regret of Squint relative to an arbitrary u is thus given by $R_T^u = \mathbb{E}_{f\sim u}\sum_{t=1}^T(\ell_t^{w_t} - \ell_t^f)$, and the variance term in (1) generalizes to $V_T^u = \mathbb{E}_{f\sim u}\sum_{t=1}^T v_t^f$ with $v_t^f = (\ell_t^{w_t} - \ell_t^f)^2$.

For such models we will use that, as shown by Koolen [2015], Squint satisfies the following quantile or 'KL' bound:

$$R_T^u \;\le\; 2\sqrt{V_T^u K_T^u} + K_T^u \tag{10}$$

which holds for every distribution u on $\mathcal{F}$ and prior π, where now $K_T^u = O(\mathrm{KL}(u\|\pi) + \ln\ln T)$ and $\mathrm{KL}(u\|\pi)$ is the KL divergence between prior π and the distribution u.

Note that (10) generalizes the countable bound (1), which is retrieved if u is taken to be a point mass on k.

Theorem 13 (Extension of Theorem 3)**.** *In any stochastic setting satisfying the (B,κ)-Bernstein Condition 1, the guarantee* (10) *for Squint implies fast rates for Squint in expectation (if there is sufficient prior mass on f that behave similarly to f^* in expectation) and with high probability (if there is sufficient prior mass on f taht are guaranteed to behave similarly to f^* on all x). That is, for all T, for any sequence $u_1, u_2, \ldots$ of distributions on $\mathcal{F}$ and sequence of numbers $C_1, C_2, \ldots$ that satisfy*

$$\mathbb{E}\left[\sum_{t=1}^T \ell_t^{u_T}\right] \le \mathbb{E}\left[\sum_{t=1}^T \ell_t^{f^*}\right] + C_T, \tag{11}$$

we have

$$\mathbb{E}[R_T^{f^*}] \;=\; O\left((K_T + C_T)^{\frac{1}{2-\kappa}} T^{\frac{1-\kappa}{2-\kappa}}\right),$$

and if (11) holds for every *sequence (x^T, y^T), then we also have for any $\delta > 0$, with probability at least $1-\delta$,*

$$R_T^{f^*} \;=\; O\left((K_T + C_T - \ln\delta)^{\frac{1}{2-\kappa}} T^{\frac{1-\kappa}{2-\kappa}}\right)$$

where $K_T := K_T^{u_T}$ from (10)*.*

While this theorem does allow us to use priors u with uncountable support, it is easiest to illustrate with priors with support on a discretized version (countable subset) of $\mathcal{F}$ which may assign probability 0 to f^*:

Example 1. Consider the classification setting where $\mathcal{J}$ is either finite or $\mathbb{N}$, and $\mathcal{F} = \bigcup_{j\in\mathcal{J}} \mathcal{F}_j$ is a finite or countable union of sub-models such that for $\delta > 0$, $\ddot{\mathcal{F}}_{j,\delta} \subset \mathcal{F}_j$ is a minimal δ-cover of $\ddot{\mathcal{F}}_j$ in the ℓ_∞-norm (that is, we require

$\sup_{f\in\mathcal{F}_j}\min_{\dot f\in\ddot{\mathcal{F}}_{j,\delta}}\sup_{x\in\mathcal{X},y\in\mathcal{Y}}\|\ell(y,f(x)-\ell(y,\dot f(x)\|\le\delta)$. Define $\Gamma := \{2^0, 2^{-1},\ldots\}$. Assume that for all j, $\mathcal{N}(\mathcal{F}_j,\delta) := |\ddot{\mathcal{F}}_{j,\delta}| < \infty$ and note that $\log\mathcal{N}(\mathcal{F}_j,\delta)$ is the metric entropy of $\mathcal{F}_j$ in the sup norm at scale δ. Let $\pi_{\mathcal{J}}$ be a probability mass function on $\mathcal{J}$ and let $\pi_{\mathbb{N}}$ be a probability distribution on $\mathbb{N}$ with $-\log\pi_{\mathcal{J}}(j)\pi_{\mathbb{N}}(k) = O(\log(jk))$ and let π be the prior on $\bigcup_{j\in\mathbf{N},\delta\in\Gamma}\ddot F_{j,\delta}$ with mass function π given by, for $f\in\ddot F_{j,2^{-k}}$, $\pi(f) = \pi_{\mathcal{J}}(j)\pi_{\mathbb{N}}(j)/\mathcal{N}(\mathcal{F}_j,2^{-k}))$. Then Theorem 13 gives the following bound in expectation (and mutatis m utandis in probability):

$$R_T^{f^*} \;=\; O\left(\left(T2^{-k} + \min_{j,k}\log(jk) + \log\mathcal{N}(\mathcal{F}_j,2^{-k})\right)^{\frac{1}{2-\kappa}} T^{\frac{1-\kappa}{2-\kappa}}\right).$$

Bounds in terms of models with bounded ℓ_∞-entropy numbers were considered before by, e.g. Gaillard and Gerchinovitz [2015] with bounded squared error loss. We note that, if $\mathcal{F}$ has logarithmic entropy numbers (e.g. $\mathcal{F} = \mathcal{F}_1$ and $\log\mathcal{N}(\mathcal{F}_1,\epsilon) = O(-\log\epsilon)$, then, by plugging in $k = \lceil\log_2 T\rceil$, we find that this cumulative regret bound is of the form $O((\log T)\cdot T^{\frac{1-\kappa}{2-\kappa}})$, the standard rate referred to in the discussion underneath Theorem 3. In the case of larger (polynomial) entropy numbers, our bounds are presumably suboptimal compared to the bounds that can be obtained by ERM, since Squint is essentially a form of an exponentially weighted forecaster that cannot exploit the chaining technique, viz. the discussion by Gaillard and Gerchinovitz [2015], Audibert [2009] and Rakhlin and Sridharan [2014]. Nevertheless, unlike ERM, Squint is robust and will continue to achieve nontrivial regret under nonstochastic, adversarially generated data, even with polynomial entropy numbers.

In practice, one may often work with $\mathcal{F}_i$ which have small (e.g. logarithmic) entropy numbers relative to the pseudo-distance $d(f_1,f_2) = \mathbb{P}(f_1(X)\neq f_2(X))$ considered by e.g. Tsybakov [2004], Audibert [2004], which may be much smaller than the ℓ_∞-numbers. In such cases, Theorem 13 can still be used to give good bounds in expectation.

Proof of Theorem 13 Consider a (for now) arbitrary sequence $u_1,u_2,\ldots$, define $K_T := K_T^{u_T}$ and $K'_T = K_T/4$, and

$$C'_T = -(R_T^{u_T} - R_T^{f^*}) \text{ or equivalently } \sum_{t=1}^T \ell_t^{u_T} = \sum_{t=1}^T \ell_t^{f^*} + C'_T.$$

One easily shows that for general $a,b,c\in\mathbb{R}$, one has $(a-b)^2/2 \le (a-c)^2+(b-c)^2$. Applying the statement with $a=\ell_t^{w_t}$, $b=\ell_t^f$ and $c=\ell^{f_t^*}$ gives $v_t^f \le 2(v_t^{f^*}+(x_t^f)^2)$. Summing over $t = 1..T$ and taking expectation over $f\sim u_T$ now gives $V_T^{u_T} \le 2V_T^{f^*} + 2E_T$ where $E_T = \mathbb{E}_{f\sim u_T}\left[\sum_{t=1}^T (x_t^f)^2\right]$.

Applying this to the bound (10) above at u_T, we get

$$R_T^{f^*} \;\le\; C'_T + 2\sqrt{(V_T^{f^*}+E_T)2K'_T} + K'_T \;=\; \inf_\eta\left\{C'_T + \eta(V_T^{f^*}+E_T) + \frac{2K'_T}{\eta} + K'_T\right\}. \tag{12}$$

We first prove an analogue to Theorem 9 for the uncountable setting, based on (12). As in that theorem, let, for given $\mathcal{F}$, $\left\{x_t^f \middle| f \in \mathcal{F}\right\}$ be the associated the excess loss family from (3), and let $\epsilon(\eta)$ be, as in (4), the corresponding maximal normalized cumulant generating function. Let K_T' be as above. Fix $\gamma \geq 0$ and let c be as in Lemma 8. Now as in the proof of Theorem 9 we have for all $f \in \mathcal{F}$, $x_t^f \in [-1,1]$ and $-x_t^f \trianglelefteq_\eta \epsilon(\eta)$ by construction of $\epsilon(\eta)$. Hence $-x_t^f \trianglelefteq_\gamma \epsilon(\gamma)$ for all $f \in \mathcal{F}$, which implies $-\mathbb{E}_{f\sim w_t} x_t^f \trianglelefteq_\gamma \epsilon(\gamma)$ and also $-\mathbb{E}_{f\sim u_T} x_t^f \trianglelefteq_\gamma \epsilon(\gamma)$, and hence by Lemma 8 (see remark below the lemma),

$$c\gamma \mathop{\mathbb{E}}_{f\sim w_t}[x_t^f]^2 - \mathop{\mathbb{E}}_{f\sim w_t}[x_t^f] \trianglelefteq_\gamma \epsilon(2\gamma)(1+c\gamma^2) \text{ and} \tag{13}$$

$$c\gamma \mathop{\mathbb{E}}_{f\sim u_T}[x_t^f]^2 - \mathop{\mathbb{E}}_{f\sim u_T}[x_t^f] \trianglelefteq_\gamma \epsilon(2\gamma)(1+c\gamma^2). \tag{14}$$

Using $r_t^{f^*} = \mathbb{E}_{f\sim w_t}\left[x_t^f\right]$, again analogously to the proof of Theorem 9, we may telescope (13) over rounds to

$$c\gamma V_T^{f^*} - R_T^{f^*} \trianglelefteq_\gamma T\epsilon(2\gamma)(1+c\gamma^2) \tag{15}$$

Now we use (12) with $\eta = \frac{c\gamma}{2}$, which implies $2R_T^{f^*} \leq 2C_T' + c\gamma(V_T^{f^*} + E_T) + 4K_T'/(c\gamma) + 2K_T'$. Combining this with (15), we find:

$$U \trianglelefteq 0 \text{ with } U = \gamma R_T^{f^*} - \left(2\gamma C_T' + c\gamma^2 E_T + \frac{4K_T'}{c} + \gamma 2K_T' + \gamma T\epsilon(2\gamma)(1+c\gamma^2)\right). \tag{16}$$

Similarly to deriving (15), using the definition of E_T, we may telescope (14) over rounds to get

$$U' \trianglelefteq 0 \text{ with } U' = c\gamma^2 E_T - \gamma C_T' - \gamma T\epsilon(2\gamma)(1+c\gamma^2). \tag{17}$$

We may now combine (16) and (17) using Lemma 7 with w a distribution that puts mass $1/2$ on random variable U and $1/2$ on U', to get $(U+U')/2 \trianglelefteq 0$, which can be rewritten to:

$$\frac{\gamma}{2}\Bigg(R_T^{f^*} - (2C_T' + c\gamma_T^E + \frac{4K_T'}{c\gamma} + 2K_T' + T\epsilon(2\gamma)(1+c\gamma^2)) \\ + c\gamma E_T - \gamma C_T' - \gamma\epsilon(2\gamma)(1+c\gamma^2)\Bigg) \trianglelefteq 0,$$

and further to

$$\frac{1}{2}R_T^{f^*} \trianglelefteq_\gamma \frac{3}{2}C_T' + \frac{K_T}{c\gamma} + T\epsilon(2\gamma)(1+c\gamma^2) + 2K_T', \tag{18}$$

which is the required analogue of the statement of Theorem 9. Note that this statement holds for *every* sequence $u_1, \ldots, u_T$, and C_T' is a random variable that depends on data (x^T, y^T).

The remainder of the proof of Theorem 13 now follows in a fashion entirely analogous to the proof of Theorem 3, as in Appendix E, where we use that we can bound C'_T by C_T, either in expectation or on all sequences; we omit further details where one uses (18) instead of the corresponding statement of Theorem 9; we omit further details.